WEST Search History

| Hide Items | Restore | Clear | Cancel |

DATE: Wednesday, April 26, 2006

Hide?	Set Name	Query	Hit Count
		DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ	
☐	L54	L53 and (table near5 space)	4
☐	L53	L52 and drop and create and sql	6
☐	L52	rdbms and (recover$3 near5 table$1) and @py<=2001	11
☐	L51	'drop table' and 'table space' and 'table id' and @py<=2000	1
☐	L50	L49 and (table space)	1
☐	L49	(drop tables) and command$1 and trigger$1 and @py<=2001	14
☐	L48	(data structure) and trigger$1 and (drop tables) and (table space) and @py<=2001	1
☐	L47	(data structure) and trigger$1 and sql and (table space) and command$1 and @py<=2001	4
☐	L46	'data structure' and 'drop table' and sql and @py<=2001	25
☐	L45	L44 and 'table space'	2
☐	L44	L43 and rdbms	38
☐	L43	'drop table' and 'create table' and sql	105
☐	L42	'drop table' and 'create table' and sql and 'table space' and @py<=2001	4
☐	L41	(database near5 management) and (table near5 space) and (drop near5 table$1) and @py<=2001	12
☐	L40	L39 and (drop near5 table$1)	3
☐	L39	(table near5 space) and rdbms and table$1 and sql and @py<=2000	38
		DB=USPT,PGPB; PLUR=YES; OP=ADJ	
☐	L38	('4945474'| '5455946'| '5758355'| '5832508'| '5884327'| '5890165')![pn]	6
		DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ	
☐	L37	6792435.pn.	2
☐	L36	L35 and (drop near5 table$1)	3
☐	L35	L34 and sql	38
☐	L34	L33 and (table near5 space$1)	117
☐	L33	(database$1 or data$base$1).ti. and @py<2000	2513003
☐	L32	(database$1 or data$base$1).ti.	2545888
☐	L31	L30 and (table space)	6
☐	L30	L29 and (drop table)	110
☐	L29	create table	3024

☐	L28	L27 and sql	1
☐	L27	L26 and (table near5 space)	12
☐	L26	(drop near5 table).clm. and @py<=2000	122
☐	L25	(table near5 id) and (drop near5 table) and (table near5 space) and (copy near5 table) and command$1 and sql and @py<=2000	0
☐	L24	dbms and sql and recover$3 and table$1 and (drop near5 table$1) and data and structure and command$1 and restore and space and history and catalog and command$1 and rollforward and @py<=2000	0
☐	L23	(recover$3 near5 table$1 and dropp$3 and relational and sql) and @py<=2000	3
☐	L22	L21 and backup	6
☐	L21	L20 and dropp$3 and creat$3 and sql	15
☐	L20	(database near5 table$1) and (table near5 space) and (data near5 structure) and attribute$1 and @py<=2000	167
☐	L19	(copy$3 near5 table$1) and (load$3 near5 table$1) and rolling and (table near5 space) and·recover$3 and flag$1 and sql and dropp$3 and creat$3 and @py<=2000	0
☐	L18	sql and command$1 and table$1 and (table near5 id$) and dropp$3 and creat$3 and modify$3 and backup and restore and (table near5 space) and attribute$1 and (data near5 structure) and @py<=2000	1
☐	L17	L16 and (dropp$3 near5 table$1)	5
☐	L16	L12 and sql	40
☐	L15	L14 and backup	2
☐	L14	L13 and recover$3	4
☐	L13	L12 and (relational near5 table$1)	36
☐	L12	(table near5 space) and (table near5 id$) and @py<=2000	745
☐	L11	L10 and attribute$1 and timestamp$3	0
☐	L10	(table$1 near5 id$) and (table near5 space$1) and (relational near5 table$1) and modify$3 and dropp$3 and creat$3 and @py<=2000	6
☐	L9	(table$1 near5 id$) and (table near5 space$1) and (relational near5 table$1) and modify$3 and dropp$3 and creat$3 and (log near5 (record$1 or file$1)) and @py<=2000	0
☐	L8	5890167.uref.	7
☐	L7	(table near5 space) and sql and command$1 and dropp$3 and creat$3 and recover$3 and @py<=2000	8
☐	L6	L5 and restor$3	2
☐	L5	(relational near5 database$1) and (dropp$3 near5 table$1) and sql and (table near5 space) and @py<=2000	8
☐	L4	L2 and sql and command$1	1
☐	L3	L2 and (dropp$3 near5 table$1)	0
☐	L2	L1 and flag$1 and recover$3	46
☐	L1	database$1 and management and table and space and dropp$3 and stor$3 and log$1 and history and restor$3 and transaction$1 and table$1 and id$ and	64

@py<=1999

END OF SEARCH HISTORY

P◉RTAL

USPTO

Search: ⦿ The ACM Digital Library   ○ The Guide

table space drop table sql

SEARCH

THE ACM DIGITAL LIBRARY

◄ Feedback Report a problem Satisfaction survey

Terms used **table space drop table sql**   Found **76,196** of **175,083**

Sort results by   [relevance ▽]   ● Save results to a Binder   Try an Advanced Search
Display results   [expanded form ▽]   ⓘ Search Tips   Try this search in The ACM Guide
                                          ☐ Open results in a new window

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next
Best 200 shown   Relevance scale ☐ ▭ ▬ ▬ ■

1 Automated SQL documentation using APL2

Rexford H. Swain
July 1991 **ACM SIGAPL APL Quote Quad , Proceedings of the international conference on APL '91 APL '91** , Volume 21 Issue 4
Publisher: ACM Press

Full text available: 📄 pdf(785.07 KB)   Additional Information: full citation, abstract, references, index terms

An application programmer working with APL2 and SQL/DS will often want to investigate and/or document the definitions of SQL tables and views, particularly when working with tables created by others. Unlike conventional APL file systems, SQL "knows" quite a bit about the objects it is storing, but this information is scattered throughout several system catalogs. An APL2 tool which combines and neatly formats available information about a table is presented.Interpretation of this information may ...

2 APL2 and SQL (tutorial session): a tutorial

Nancy Wheeler
August 1989 **Proceedings of the ACM/SIGAPL conference on APL as a tool of thought (session tutorials)**
Publisher: ACM Press

Full text available: 📄 pdf(1.97 MB)   Additional Information: full citation, references, index terms

3 SQLEM: fast clustering in SQL using the EM algorithm

Carlos Ordonez, Paul Cereghini
May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data SIGMOD '00** , Volume 29 Issue 2
Publisher: ACM Press

Full text available: 📄 pdf(1.07 MB)   Additional Information: full citation, abstract, references, citings, index terms

Clustering is one of the most important tasks performed in Data Mining applications. This paper presents an efficient SQL implementation of the EM algorithm to perform clustering in very large databases. Our version can effectively handle high dimensional data, a high number of clusters and more importantly, a very large number of data records. We present three strategies to implement EM in SQL: horizontal, vertical and a hybrid one. We expect this work to be useful for data mining programmer ...

4

An extension of the database language SQL to capture more relational concepts

Gottfried Vossen, Jim Yacabucci
November 1988 **ACM SIGMOD Record**, Volume 17 Issue 4
Publisher: ACM Press
Full text available: pdf(623.92 KB)   Additional Information: full citation, abstract, index terms

An extension of the database language SQL is described which introduces several new concepts into the language that are standard in the relational model, but surprisingly not present in available SQL-based systems such as SQL/DS and DB2.

5 R* optimizer validation and performance evaluation for local queries

Lothar F. Mackert, Guy M. Lohman
June 1986 **ACM SIGMOD Record , Proceedings of the 1986 ACM SIGMOD international conference on Management of data SIGMOD '86**, Volume 15 Issue 2
Publisher: ACM Press

Full text available: pdf(1.43 MB)   Additional Information: full citation, abstract, references, citings, index terms

Few database query optimizer models have been validated against actual performance. This paper presents the methodology and results of a thorough validation of the optimizer and evaluation of the performance of the experimental distributed relational database management system R*, which inherited and extended to a distributed environment the optimization algorithms of System R. Optimizer estimated costs and actual R* resources consumed were written ...

6 Searching in metric spaces with user-defined and approximate distances

Paolo Ciaccia, Marco Patella
December 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 4
Publisher: ACM Press

Full text available: pdf(555.89 KB)   Additional Information: full citation, abstract, references, citings, index terms

Novel database applications, such as multimedia, data mining, e-commerce, and many others, make intensive use of similarity queries in order to retrieve the objects that better fit a user request. Since the effectiveness of such queries improves when the user is allowed to personalize the similarity criterion according to which database objects are evaluated and ranked, the development of access methods able to efficiently support user-defined similarity queries becomes a basic requirement. In t ...

Keywords: Distance metrics, user-defined queries

7 Research sessions: stream management: Holistic UDAFs at streaming speeds

Graham Cormode, Theodore Johnson, Flip Korn, S. Muthukrishnan, Oliver Spatscheck, Divesh Srivastava
June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**
Publisher: ACM Press
Full text available: pdf(264.85 KB)   Additional Information: full citation, abstract, references, citings

Many algorithms have been proposed to approximate holistic aggregates, such as quantiles and heavy hitters, over data streams. However, little work has been done to explore what techniques are required to incorporate these algorithms in a data stream query processor, and to make them useful in practice.In this paper, we study the performance implications of using user-defined aggregate functions (UDAFs) to incorporate selection-based and sketch-based algorithms for holistic aggregates into a dat ...

8 Intrusion detection and prevention: On deriving unknown vulnerabilities from zero-day ▣
polymorphic and metamorphic worm exploits
Jedidiah R. Crandall, Zhendong Su, S. Felix Wu, Frederic T. Chong
November 2005 **Proceedings of the 12th ACM conference on Computer and
communications security CCS '05**
Publisher: ACM Press
Full text available: 📄 pdf(334.95 KB)   Additional Information: full citation, abstract, references, index terms

Vulnerabilities that allow worms to hijack the control flow of each host that they spread to
are typically discovered months before the worm outbreak, but are also typically
discovered by third party researchers. A determined attacker could discover vulnerabilities
as easily and create zero-day worms for vulnerabilities unknown to network defenses. It is
important for an analysis tool to be able to generalize from a new exploit observed and
derive protection for the vulnerability.Many researcher ...

Keywords: honeypots, metamorphism, polymorphic worms, polymorphism, symbolic
execution, worms

9 FS2: dynamic data replication in free disk space for improving disk performance and ▣
energy consumption
Hai Huang, Wanda Hung, Kang G. Shin
October 2005 **ACM SIGOPS Operating Systems Review , Proceedings of the twentieth
ACM symposium on Operating systems principles SOSP '05**, Volume 39 Issue
5
Publisher: ACM Press
Full text available: 📄 pdf(542.63 KB)   Additional Information: full citation, abstract, references, index terms

Disk performance is increasingly limited by its head positioning latencies, i.e., seek time
and rotational delay. To reduce the head positioning latencies, we propose a novel
technique that *dynamically* places copies of data in file system's *free blocks* according to
the disk access patterns observed at runtime. As one or more replicas can now be
accessed in addition to their original data block, choosing the "nearest" replica that
provides fastest access can significantly improve pe ...

Keywords: data replication, disk layout reorganization, dynamic file system, free disk
space

10 Exploiting early sorting and early partitioning for decision support query processing ▣
J. Claussen, A. Kemper, D. Kossmann, C. Wiesner
December 2000 **The VLDB Journal — The International Journal on Very Large Data
Bases**, Volume 9 Issue 3
Publisher: Springer-Verlag New York, Inc.
Full text available: 📄 pdf(478.23 KB)   Additional Information: full citation, abstract, index terms

Decision support queries typically involve several joins, a grouping with aggregation,
and/or sorting of the result tuples. We propose two new classes of query evaluation
algorithms that can be used to speed up the execution of such queries. The algorithms
are based on (1) *early sorting* and (2) *early partitioning* or a combination of both. The
idea is to push the sorting and/or the partitioning to the leaves, i.e., the base relations, of
the query evaluation plans (QEPs) and ...

Keywords: Decision Support Systems, Early sorting and partitioning, Hash joins and
hash teams, Performance evaluation, Query processing and optimization

11 Using transact-SQL and simulation techniques to create virtual M&M'S
Robin M. Snyder
December 2002 **Journal of Computing Sciences in Colleges**, Volume 18 Issue 2
Publisher: Consortium for Computing Sciences in Colleges
Full text available: pdf(70.11 KB)   Additional Information: full citation, abstract, references, index terms

> The M&M problem is an example of a nontrivial yet simple method for summarizing
> sampled data that are then used in making decisions. Previous work by the author used
> client-side JavaScript and Active Server Pages to create virtual M&M's. This paper moves
> the generation of the virtual M&M's to a SQL Server database server. In doing so, a
> number of interesting and practical problems in database design and implementation are
> explored. In particular, the use of the programming language Transact-SQL ...

12 Databases: SQL DOM: compile time checking of dynamic SQL statements
Russell A. McClure, Ingolf H. Krüger
May 2005 **Proceedings of the 27th international conference on Software engineering**
Publisher: ACM Press
Full text available: pdf(353.48 KB)   Additional Information: full citation, abstract, references, index terms

> Most object oriented applications that involve persistent data interact with a relational
> database. The most common interaction mechanism is a call level interface (CLI) such as
> ODBC or JDBC. While there are many advantages to using a CLI -- expressive power and
> performance being two of the most key -- there are also drawbacks. Applications
> communicate through a CLI by constructing strings that contain SQL statements. These
> SQL statements are only checked for correctness at runtime, tend to be f ...
>
> **Keywords**: SQL, SQL DOM, SQL injection, SQL strings, dynamic SQL, impedance
> mismatch

13 Object oriented relational database with SQL interface
Behrooz Seyed-Abbassi
March 1993 **Proceedings of the 1993 ACM conference on Computer science**
Publisher: ACM Press
Full text available: pdf(1.43 MB)   Additional Information: full citation, abstract, references, index terms

> An object oriented relational database management system to support heterogeneous
> object classes of statistical (numeric and text), image, text and sound information is
> considered. By employing a user friendly interface and Structured Query Language (SQL)
> capability at the user level, this database with a three level architecture utilizes an
> extended SQL to operate on complex objects and to support database processing
> functions such as retrieve, join and overlay. These operations also suppo ...

14 Research sessions: consistency and availability: Relaxed currency and consistency:
how to say "good enough" in SQL
Hongfei Guo, Per-Åke Larson, Raghu Ramakrishnan, Jonathan Goldstein
June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on
Management of data**
Publisher: ACM Press
Full text available: pdf(606.81 KB)   Additional Information: full citation, abstract, references, citings

> Despite the widespread and growing use of asynchronous copies to improve scalability,
> performance and availability, this practice still lacks a firm semantic foundation.
> Applications are written with some understanding of which queries can use data that is
> not entirely current and which copies are "good enough"; however, there are neither

explicit requirements nor guarantees. We propose to make this knowledge available to the DBMS through explicit currency and consistency (C&C) constraints in qu ...

15 A critique of ANSI SQL isolation levels

Hal Berenson, Phil Bernstein, Jim Gray, Jim Melton, Elizabeth O'Neil, Patrick O'Neil
May 1995 **ACM SIGMOD Record , Proceedings of the 1995 ACM SIGMOD international conference on Management of data SIGMOD '95**, Volume 24 Issue 2
Publisher: ACM Press

Full text available: pdf(1.20 MB)   Additional Information: full citation, abstract, references, citings, index terms

ANSI SQL-92 [MS, ANSI] defines Isolation *Levels* in terms of *phenomena:* Dirty Reads, Non-Repeatable Reads, and Phantoms. This paper shows that these phenomena and the ANSI SQL definitions fail to properly characterize several popular isolation levels, including the standard locking implementations of the levels covered. Ambiguity in the statement of the phenomena is investigated and a more formal statement is arrived at; in addition new phenomena that better characterize isolation t ...

16 Research sessions: indexing and tuning: Integrating vertical and horizontal partitioning into automated physical database design

Sanjay Agrawal, Vivek Narasayya, Beverly Yang
June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**
Publisher: ACM Press

Full text available: pdf(181.69 KB)   Additional Information: full citation, abstract, references

In addition to indexes and materialized views, horizontal and vertical partitioning are important aspects of physical design in a relational database system that significantly impact performance. Horizontal partitioning also provides manageability; database administrators often require indexes and their underlying tables partitioned identically so as to make common operations such as backup/restore easier. While partitioning is important, incorporating partitioning makes the problem of automatin ...

17 Using Applications of Data Versioning in Database Application Development

Ramkrishna Chatterjee, Gopalan Arun, Sanjay Agarwal, Ben Speckhard, Ramesh Vasudevan
May 2004 **Proceedings of the 26th International Conference on Software Engineering ICSE '04**
Publisher: IEEE Computer Society

Full text available: pdf(166.57 KB)   Additional Information: full citation, abstract, citings, index terms

Database applications such as enterprise resource planning systems and customer relationship management systems are widely used software systems. Development and testing of database applications is difficult because the program execution depends on the persistent state stored in the database. In this paper we show that how versioning of the persistent data stored in the database can solve some critical problems in the development and testing of database applications can be solved by vers ...

18 At the Forge: Speaking SQL

Reuven Lerner
September 1997 **Linux Journal**
Publisher: Specialized Systems Consultants, Inc.

Full text available: html(20.68 KB)   Additional Information: full citation, citings, index terms

19

Querying web metadata: Native score management and text support in databases

Gültekin Özsoyo☐lu, Ismail Sengör Altingövde, Abdullah Al-Hamdani, Selma Ayşe Özel,
Özgür Ulusoy, Zehra Meral özsoyo☐lu
December 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 4
Publisher: ACM Press
Full text available: pdf(737.76 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>index terms</u>

> In this article, we discuss the issues involved in adding a native score management
> system to object-relational databases, to be used in querying Web metadata (that
> describes the semantic content of Web resources). The Web metadata model is based on
> topics (representing entities), relationships among topics (called *metalinks*), and
> importance scores (sideway values) of topics and metalinks. We extend database
> relations with scoring functions and importance scores. We add to SQL score-manag ...

> **Keywords**: Score management for Web applications

20 <u>Loading databases using dataflow parallelism</u>   ■

Tom Barclay, Robert Barnes, Jim Gray, Prakash Sundaresan
December 1994 **ACM SIGMOD Record**, Volume 23 Issue 4
Publisher: ACM Press
Full text available: pdf(1.49 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>citings</u>, <u>index terms</u>

> This paper describes a parallel database load prototype for Digital's Rdb database product.
> The prototype takes a dataflow approach to database parallelism. It includes an *explorer*
> that discovers and records the cluster configuration in a database, a *client* CUI interface
> that gathers the load job description from the user and from the Rdb catalogs, and an
> *optimizer* that picks the best parallel execution plan and records it in a *web* data
> structure. The web describes th ...

Results 1 - 20 of 200   Result page: **1** <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u>   <u>next</u>

Useful downloads: <u>Adobe Acrobat</u>   <u>QuickTime</u>   <u>Windows Media Player</u>   <u>Real Player</u>

P⊛RTAL
USPTO

THE ACM DIGITAL LIBRARY

Terms used **recovering** **rollback** **table** **space** **drop** **table** **sql**   Found **52,963** of **175,083**

Sort results by: [relevance ▽]

Display results: [expanded form ▽]

❖ Save results to a Binder

▣ Search Tips

☐ Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next
Best 200 shown   Relevance scale ☐ ▭ ▬ ▬ ■

1 Garbage collection for a client-server persistent object store   ■

Laurent Amsaleg, Michael J. Franklin, Olivier Gruber
August 1999 **ACM Transactions on Computer Systems (TOCS)**, Volume 17 Issue 3
Publisher: ACM Press

Full text available: ▨ pdf(267.18 KB)   Additional Information: full citation, abstract, references, citings, index terms, review

We describe an efficient server-based algorithm for garbage collecting persistent object stores in a client-server environmnet. The algorithm is incremental and runs concurrently with client transactions. Unlike previous algorithms, it does not hold any transactional locks on data and does non require callbacks to clients. It is fault-tolerant, but performs very little logging. The algorithm has been designed to be integrated into existing systems, and therefore it works with standard i ...

Keywords: client-server system, logging, persistent object-store, recovery

2 ARIES: a transaction recovery method supporting fine-granularity locking and partial   ■
rollbacks using write-ahead logging
C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz
March 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 1
Publisher: ACM Press

Full text available: ▨ pdf(5.23 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

Keywords: buffer management, latching, locking, space management, write-ahead logging

3 Computing curricula 2001   ■
September 2001 **Journal on Educational Resources in Computing (JERIC)**

Publisher: ACM Press
Full text available: 📄 pdf(613.63 KB)   Additional Information: <u>full citation</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>
             🔲 html(2.78 KB)

4 <u>Final report of the ANSI/X3/SPARC DBS-SG relational database task group</u>   ⬛
July 1982 **ACM SIGMOD Record**, Volume 12 Issue 4
Publisher: ACM Press
Full text available: 📄 pdf(4.69 MB)   Additional Information: <u>full citation</u>

5 <u>Performance analysis of recovery techniques</u>   ⬛
Andreas Reuter
December 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 4
Publisher: ACM Press

Full text available: 📄 pdf(2.47 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>, <u>review</u>

Various logging and recovery techniques for centralized transaction-oriented database systems under performance aspects are described and discussed. The classification of functional principles that has been developed in a companion paper is used as a terminological basis. In the main sections, a set of analytic models is introduced and evaluated in order to compare the performance characteristics of nine different recovery techniques with respect to four key parameters and a set of other pa ...

6 <u>Anatomy of a native XML base management system</u>   ⬛
T. Fiebig, S. Helmer, C.-C. Kanne, G. Moerkotte, J. Neumann, R. Schiele, T. Westmann
December 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue 4
Publisher: Springer-Verlag New York, Inc.
Full text available: 📄 pdf(300.97 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>citings</u>, <u>index terms</u>

Several alternatives to manage large XML document collections exist, ranging from file systems over relational or other database systems to specifically tailored XML base management systems. In this paper we give a tour of Natix, a database management system designed from scratch for storing and processing XML data. Contrary to the common belief that management of XML data is just another application for traditional databases like relational systems, we illustrate how almost every component in a ...

Keywords: Database, XML

7 <u>Towards effective and efficient free space management</u>   ⬛
Mark L. McAuliffe, Michael J. Carey, Marvin H. Solomon
June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96**, Volume 25 Issue 2
Publisher: ACM Press

Full text available: 📄 pdf(1.34 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

An important problem faced by many database management systems is the "online object placement problem"--the problem of choosing a disk page to hold a newly allocated object. In the absence of clustering criteria, the goal is to maximize storage utilization. For main-memory based systems, simple heuristics exist that provide reasonable space utilization in the worst case and excellent utilization in typical cases. However, the storage

management problem for databases includes significant additio ...

8 <u>Integrating reliable memory in databases</u>

Wee Teck Ng, Peter M. Chen

August 1998 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 7 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: <u>pdf(123.18 KB)</u>   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>index terms</u>

Recent results in the Rio project at the University of Michigan show that it is possible to create an area of main memory that is as safe as disk from operating system crashes. This paper explores how to integrate the reliable memory provided by the Rio file cache into a database system. Prior studies have analyzed the performance benefits of reliable memory; we focus instead on how different designs affect reliability. We propose three designs for integrating reliable memory into databases: non ...

Keywords: Main memory database system (MMDB), Recovery, Reliability

9 <u>Rx: treating bugs as allergies---a safe method to survive software failures</u>

Feng Qin, Joseph Tucek, Jagadeesan Sundaresan, Yuanyuan Zhou

October 2005 **ACM SIGOPS Operating Systems Review , Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05**, Volume 39 Issue 5

Publisher: ACM Press

Full text available: <u>pdf(245.29 KB)</u>   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>index terms</u>

Many applications demand availability. Unfortunately, software failures greatly reduce system availability. Prior work on surviving software failures suffers from one or more of the following limitations: Required application restructuring, inability to address deterministic software bugs, unsafe speculation on program execution, and long recovery time.This paper proposes an innovative *safe* technique, called Rx, which can quickly recover programs from many types of software bugs, both det ...

Keywords: availability, bug, reliability, software failure

10 <u>APL2 and SQL (tutorial session): a tutorial</u>

Nancy Wheeler

August 1989 **Proceedings of the ACM/SIGAPL conference on APL as a tool of thought (session tutorials)**

Publisher: ACM Press

Full text available: <u>pdf(1.97 MB)</u>   Additional Information: <u>full citation</u>, <u>references</u>, <u>index terms</u>

11 <u>Intrusion detection and prevention: Automatic diagnosis and response to memory corruption vulnerabilities</u>

Jun Xu, Peng Ning, Chongkyung Kil, Yan Zhai, Chris Bookholt

November 2005 **Proceedings of the 12th ACM conference on Computer and communications security CCS '05**

Publisher: ACM Press

Full text available: <u>pdf(285.20 KB)</u>   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>index terms</u>

Cyber attacks against networked computers have become relentless in recent years. The most common attack method is to exploit memory corruption vulnerabilities such as buffer overflow and format string bugs. This paper presents a technique to automatically

identify both known and unknown memory corruption vulnerabilities. Based on the observation that a randomized program usually crashes upon a memory corruption attack, this technique uses the crash as a trigger to initiate an automatic diagnosi ...

Keywords: attack diagnosis, memory corruption attack, message filtering, randomization

12 Vigilante: end-to-end containment of internet worms

Manuel Costa, Jon Crowcroft, Miguel Castro, Antony Rowstron, Lidong Zhou, Lintao Zhang, Paul Barham
October 2005 **ACM SIGOPS Operating Systems Review , Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05**, Volume 39 Issue 5
Publisher: ACM Press
Full text available: pdf(329.29 KB)   Additional Information: full citation, abstract, references, index terms

Worm containment must be automatic because worms can spread too fast for humans to respond. Recent work has proposed network-level techniques to automate worm containment; these techniques have limitations because there is no information about the vulnerabilities exploited by worms at the network level. We propose Vigilante, a new end-to-end approach to contain worms automatically that addresses these limitations. Vigilante relies on collaborative worm detection at end hosts, but does not requir ...

Keywords: control flow analysis, data flow analysis, self-certifying alerts, worm containment

13 Consistency and replication: Decentralized weighted voting for P2P data management

Maya Rodrig, Anthony LaMarca
September 2003 **Proceedings of the 3rd ACM international workshop on Data engineering for wireless and mobile access**
Publisher: ACM Press
Full text available: pdf(448.69 KB)   Additional Information: full citation, abstract, references, citings, index terms

This paper presents a decentralized variant of David Gifford's classic weighted-voting scheme for managing replicated data. Weighted voting offers a familiar consistency model and supports on-line replica reconfiguration. These properties make it a good fit for applications in the pervasive computing domain. By distributing versioned metadata along with data replicas, and managing access to both data and metadata with the same quorums, our algorithm supports a peer-to-peer environment with dynam ...

Keywords: consistency, peer-to-peer data management, pervasive computing, quorums, weighted voting

14 The O2 system

O. Deux
October 1991 **Communications of the ACM**, Volume 34 Issue 10
Publisher: ACM Press
Full text available: pdf(7.18 MB)   Additional Information: full citation, references, citings, index terms

Keywords: O2, Object-oriented database systems

15 Industrial sessions: beyond relational tables: Coordinating backup/recovery and data consistency between database and file systems

Suparna Bhattacharya, C. Mohan, Karen W. Brannon, Inderpal Narang, Hui-I Hsiao, Mahadevan Subramanian

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available: pdf(1.44 MB)   Additional Information: full citation, abstract, references, index terms

Managing a combined store consisting of database data and file data in a robust and consistent manner is a challenge for database systems and content management systems. In such a hybrid system, images, videos, engineering drawings, etc. are stored as files on a file server while meta-data referencing/indexing such files is created and stored in a relational database to take advantage of efficient search. In this paper we describe solutions for two potentially problematic aspects of such a data ...

Keywords: DB2, content management, database backup, database recovery, datalinks

16 Research session: database architectures for new hardware: Optimistic intra-transaction parallelism on chip multiprocessors

Christopher B. Colohan, Anastassia Ailamaki, J. Gregory Steffan, Todd C. Mowry

August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05**

Publisher: VLDB Endowment

Full text available: pdf(418.30 KB)   Additional Information: full citation, abstract, references, index terms

With the advent of chip multiprocessors, exploiting intra-transaction parallelism is an attractive way of improving transaction performance. However, exploiting intra-transaction parallelism in existing database systems is difficult, for two reasons: first, significant changes are required to avoid races or conflicts within the DBMS, and second, adding threads to transactions requires a high level of sophistication from transaction programmers. In this paper we show how dividing a transaction in ...

17 A prototype implementation of the SQL Ada module extension (SAME) method

Allison LeClair, Susan Phillips

December 1990 **Proceedings of the conference on TRI-ADA '90**

Publisher: ACM Press

Full text available: pdf(1.20 MB)   Additional Information: full citation, abstract, references, citings

As Ada becomes more widespread, the ability to access commercial database technologies through Ada systems becomes a significant issue. Researchers throughout our industry are investigating interface approaches between Ada and these technologies, including language bindings between Ada and SQL, a relational data base language. This paper presents a recent implementation of one such binding—the SQL Ada Module Extension (SAME) method.

18 Automated SQL documentation using APL2

Rexford H. Swain

July 1991 **ACM SIGAPL APL Quote Quad , Proceedings of the international conference on APL '91 APL '91**, Volume 21 Issue 4

Publisher: ACM Press

Full text available: pdf(785.07 KB)   Additional Information: full citation, abstract, references, index terms

An application programmer working with APL2 and SQL/DS will often want to investigate

and/or document the definitions of SQL tables and views, particularly when working with tables created by others. Unlike conventional APL file systems, SQL "knows" quite a bit about the objects it is storing, but this information is scattered throughout several system catalogs. An APL2 tool which combines and neatly formats available information about a table is presented.Interpretation of this information may ...

19 Client-server computing in mobile environments

Jin Jing, Abdelsalam Sumi Helal, Ahmed Elmagarmid
June 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 2
Publisher: ACM Press

Full text available: pdf(233.31 KB)   Additional Information: full citation, abstract, references, citings, index terms, review

Recent advances in wireless data networking and portable information appliances have engendered a new paradigm of computing, called mobile computing, in which users carrying portable devices have access to data and information services regardless of their physical location or movement behavior. In the meantime, research addressing information access in mobile environments has proliferated. In this survey, we provide a concrete framework and categorization of the various way ...

Keywords: application adaptation, cache invalidation, caching, client/server, data dissemination, disconnected operation, mobile applications, mobile client/server, mobile compuing, mobile data, mobility awareness, survey, system application

20 Managing update conflicts in Bayou, a weakly connected replicated storage system

D. B. Terry, M. M. Theimer, Karin Petersen, A. J. Demers, M. J. Spreitzer, C. H. Hauser
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles SOSP '95**, Volume 29 Issue 5
Publisher: ACM Press

Full text available: pdf(1.56 MB)   Additional Information: full citation, references, citings, index terms

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10 next

PORTAL
USPTO

Search: ⊙ The ACM Digital Library   ○ The Guide

table space databases   SEARCH

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used **table space databases**   Found **60,735** of **175,083**

Sort results by   [relevance ▽]   ◆ Save results to a Binder   Try an Advanced Search
Display results   [expanded form ▽]   ☑ Search Tips   Try this search in The ACM Guide
                                       ☐ Open results in a new window

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next
Best 200 shown   Relevance scale ☐ ☐ ■ ■ ■

1 Remotely-sensed geophysical databases: experience and implications for   ■
   generalized DBMS
   Guy M. Lohman, Joseph C. Stoltzfus, Anita N. Benson, Michael D. Martin, Alfonso F. Cardenas
   May 1983 **ACM SIGMOD Record , Proceedings of the 1983 ACM SIGMOD international
            conference on Management of data SIGMOD '83**, Volume 13 Issue 4
   Publisher: ACM Press
   Full text available: 📄 pdf(1.85 MB)   Additional Information: full citation, abstract, references, citings

   This paper presents the characteristics of scientific remotely-sensed databases that are
   relevant to --- and pose unique challenges for --- general-purpose database management
   systems (DBMSs). We describe a prototype system that integrates geophysical data and
   its metadata from both satellite and *in situ* sources, using a relational general-purpose
   DBMS to manage the catalog and observational data, and a video optical disk to archive
   images. Based upon our experience with this application, ...

2 I/O reference behavior of production database workloads and the TPC benchmarks— ■
   an analysis at the logical level
   Windsor W. Hsu, Alan Jay Smith, Honesty C. Young
   March 2001 **ACM Transactions on Database Systems (TODS)**, Volume 26 Issue 1
   Publisher: ACM Press
   Full text available: 📄 pdf(5.42 MB)   Additional Information: full citation, abstract, references, citings, index
                                                                 terms

   As improvements in processor performance continue to far outpace improvements in
   storage performance, I/O is increasingly the bottleneck in computer systems, especially in
   large database systems that manage huge amoungs of data. The key to achieving good
   I/O performance is to thoroughly understand its characteristics. In this article we present
   a comprehensive analysis of the logical I/O reference behavior of the peak
   productiondatabase workloads from ten of the world's largest corporatio ...

   Keywords: I/O, TPC benchmarks, caching, locality, prefetching, production database
   workloads, reference behavior, sequentiality, workload characterization

3 Hashing by proximity to process duplicates in spatial databases   ■
   Walid G. Aref, Hanan Samet
   November 1994 **Proceedings of the third international conference on Information and
                knowledge management**

Publisher: ACM Press

Full text available: pdf(952.84 KB)   Additional Information: full citation, abstract, references, citings, index terms

In a spatial database, an object may extend arbitrarily in space. As a result, many spatial data structures (e.g., the quadtree, the cell tree, the R+-tree) represent an object by partitioning it into multiple, yet simple, pieces, each of which is stored separately inside the data structure. Many operations on these data structures are likely to produce duplicate results because of the multiplicity of object pieces. A novel approach for duplicate processing based on pro ...

4 A new page table for 64-bit address spaces

M. Talluri, M. D. Hill, Y. A. Khalidi
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles SOSP '95**, Volume 29 Issue 5
Publisher: ACM Press
Full text available: pdf(1.97 MB)   Additional Information: full citation, references, citings, index terms

5 Packet classification using tuple space search

V. Srinivasan, S. Suri, G. Varghese
August 1999 **ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication SIGCOMM '99**, Volume 29 Issue 4
Publisher: ACM Press

Full text available: pdf(1.46 MB)   Additional Information: full citation, abstract, references, citings, index terms

Routers must perform packet classification at high speeds to efficiently implement functions such as firewalls and QoS routing. Packet classification requires matching each packet against a database of filters (or rules), and forwarding the packet according to the highest priority filter. Existing filter schemes with fast lookup time do not scale to large filter databases. Other more scalable schemes work for 2-dimensional filters, but their lookup times degrade quickly with each additional dime ...

6 Research track papers: Scalable mining of large disk-based graph databases

Chen Wang, Wei Wang, Jian Pei, Yongtai Zhu, Baile Shi
August 2004 **Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining KDD '04**
Publisher: ACM Press
Full text available: pdf(214.28 KB)   Additional Information: full citation, abstract, references, index terms

Mining frequent structural patterns from graph databases is an interesting problem with broad applications. Most of the previous studies focus on pruning unfruitful search subspaces effectively, but few of them address the mining on large, disk-based databases. As many graph databases in applications cannot be held into main memory, scalable mining of large, disk-based graph databases remains a challenging problem. In this paper, we develop an effective index structure, *ADI* (for <u> ...

Keywords: frequent graph pattern, graph database, graph mining, index

7 Storing and querying XML data using denormalized relational databases

Andrey Balmin, Yannis Papakonstantinou
March 2005 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 14 Issue 1

Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(397.97 KB)   Additional Information: full citation, abstract

XML database systems emerge as a result of the acceptance of the XML data model. Recent works have followed the promising approach of building XML database management systems on underlying RDBMS's. Achieving query processing performance reduces to two questions: (i) How should the XML data be decomposed into data that are stored in the RDBMS? (ii) How should the XML query be translated into an efficient plan that sends one or more SQL queries to the underlying RDBMS and combines the data ...

8 Near perfect hash table for image databases   ■
Sanjiv K. Bhatia, Chaman L. Sabharwal
February 1996 **Proceedings of the 1996 ACM symposium on Applied Computing**
Publisher: ACM Press
Full text available: pdf(572.57 KB)   Additional Information: full citation, references, index terms

Keywords: 2D string, hash table, image database indexing

9 Searching in metric spaces with user-defined and approximate distances   ■
Paolo Ciaccia, Marco Patella
December 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 4
Publisher: ACM Press
Full text available: pdf(555.89 KB)   Additional Information: full citation, abstract, references, citings, index terms

Novel database applications, such as multimedia, data mining, e-commerce, and many others, make intensive use of similarity queries in order to retrieve the objects that better fit a user request. Since the effectiveness of such queries improves when the user is allowed to personalize the similarity criterion according to which database objects are evaluated and ranked, the development of access methods able to efficiently support user-defined similarity queries becomes a basic requirement. In t ...

Keywords: Distance metrics, user-defined queries

10 Spatio-temporal database support for legacy applications   ■
Michael Böhlen, Christian S. Jensen, Bjørn Skjellaug
February 1998 **Proceedings of the 1998 ACM symposium on Applied Computing**
Publisher: ACM Press
Full text available: pdf(955.80 KB)   Additional Information: full citation, references, citings, index terms

Keywords: SQL, language design, legacy software, multi-dimensional data, spatial, spatio-temporal data, temporal

11 ARIES: a transaction recovery method supporting fine-granularity locking and partial ■
rollbacks using write-ahead logging
C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz
March 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 1
Publisher: ACM Press
Full text available: Additional Information: full citation, abstract, references, citings, index

         pdf(5.23 MB)   terms, review

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

Keywords: buffer management, latching, locking, space management, write-ahead logging

12 A universal-scheme approach to statistical databases containing homogeneous summary tables
Francesco M. Malvestuto
December 1993 **ACM Transactions on Database Systems (TODS)**, Volume 18 Issue 4
Publisher: ACM Press

Full text available: pdf(2.00 MB)   Additional Information: full citation, references, citings, index terms, review

Keywords: bipartite graph, category relation, query-answering system, statistical database, summary table, universal classification scheme

13 Empirical performance evaluation of concurrency and coherency control protocols for database sharing systems
Erhard Rahm
June 1993 **ACM Transactions on Database Systems (TODS)**, Volume 18 Issue 2
Publisher: ACM Press

Full text available: pdf(3.37 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

Database Sharing (DB-sharing) refers to a general approach for building a distributed high performance transaction system. The nodes of a DB-sharing system are locally coupled via a high-speed interconnect and share a common database at the disk level. This is also known as a "shared disk" approach. We compare database sharing with the database partitioning (shared nothing) approach and discuss the functional DBMS components that require new and coordinated solutions for DB-shar ...

Keywords: coherency control, concurrency control, database partitioning, database sharing, performance analysis, shared disk, shared nothing, trace-driven simulation

14 XML query processing I: Dynamic sample selection for approximate query processing
Brian Babcock, Surajit Chaudhuri, Gautam Das
June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**
Publisher: ACM Press

Full text available: pdf(260.80 KB)   Additional Information: full citation, abstract, references, citings, index terms

In decision support applications, the ability to provide fast approximate answers to aggregation queries is desirable. One commonly-used technique for approximate query answering is sampling. For many aggregation queries, appropriately constructed biased

(non-uniform) samples can provide more accurate approximations than a uniform sample. The optimal type of bias, however, varies from query to query. In this paper, we describe an approximate query processing technique that dynamically constructs ...

15 DB-IR-1 (databases and information retrieval): indexing and query processing effiency: Indexing text data under space constraints
Bijit Hore, Hakan Hacigumus, Bala Iyer, Sharad Mehrotra
November 2004 **Proceedings of the thirteenth ACM international conference on Information and knowledge management CIKM '04**
Publisher: ACM Press
Full text available: pdf(673.79 KB)   Additional Information: full citation, abstract, references, index terms

An important class of queries is the LIKE predicate in SQL. In the absence of an index, LIKE queries are subject to performance degradation. The notion of indexing on substrings (or <i>q</i>-grams) has been explored earlier without sufficient consideration of efficiency. <i>q</i>-grams are used to prune away rows that do not qualify for the query. The problem is to identify a finite number of grams subject to storage constraint that gives maximal pruning for a given query work ...

Keywords: B-tree, SQL, index, like queries, q-grams

16 Scalable high-speed prefix matching
Marcel Waldvogel, George Varghese, Jon Turner, Bernhard Plattner
November 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 4
Publisher: ACM Press

Full text available: pdf(933.02 KB)   Additional Information: full citation, abstract, references, citings, index terms

Finding the longest matching prefix from a database of keywords is an old problem with a number of applications, ranging from dictionary searches to advanced memory management to computational geometry. But perhaps today's most frequent best matching prefix lookups occur in the Internet, when forwarding packets from router to router. Internet traffic volume and link speeds are rapidly increasing; at the same time, a growing user population is increasing the size of routing tables against which p ...

Keywords: collision resolution, forwarding lookups, high-speed networking

17 Office-by-example: an integrated office system and database manager
Kyu-Young Whang, Art Ammann, Anthony Bolmarcich, Maria Hanrahan, Guy Hochgesang, Kuan-Tsae Huang, Al Khorasani, Ravi Krishnamurthy, Gary Sockut, Paula Sweeney, Vance Waddle, Moshé Zloof
October 1987 **ACM Transactions on Information Systems (TOIS)**, Volume 5 Issue 4
Publisher: ACM Press

Full text available: pdf(2.86 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

Office-by-Example (OBE) is an integrated office information system that has been under development at IBM Research. OBE, an extension of Query-by-Example, supports various office features such as database tables, word processing, electronic mail, graphics, images, and so forth. These seemingly heterogeneous features are integrated through a language feature called example elements. Applications involving example elements are processed by the database manager, an integrated ...

18

The Logical Record Access Approach to Database Design

Toby J. Teorey, James P. Fry
June 1980 **ACM Computing Surveys (CSUR)**, Volume 12 Issue 2
Publisher: ACM Press
Full text available: pdf(2.81 MB)   Additional Information: full citation, references, citings, index terms

19 XSB as an efficient deductive database engine   ■

Konstantinos Sagonas, Terrance Swift, David S. Warren
May 1994 **ACM SIGMOD Record , Proceedings of the 1994 ACM SIGMOD international conference on Management of data SIGMOD '94**, Volume 23 Issue 2
Publisher: ACM Press

Full text available: pdf(1.24 MB)   Additional Information: full citation, abstract, references, citings, index terms

This paper describes the XSB system, and its use as an in-memory deductive database engine. XSB began from a Prolog foundation, and traditional Prolog systems are known to have serious deficiencies when used as database systems. Accordingly, XSB has a fundamental bottom-up extension, introduced through tabling (or memoing)[4], which makes it appropriate as an underlying query engine for deductive database systems. Because it eliminates redundant computation, the tabling extension makes XSB ...

20 Using Applications of Data Versioning in Database Application Development   ■

Ramkrishna Chatterjee, Gopalan Arun, Sanjay Agarwal, Ben Speckhard, Ramesh Vasudevan
May 2004 **Proceedings of the 26th International Conference on Software Engineering ICSE '04**
Publisher: IEEE Computer Society
Full text available: pdf(166.57 KB)   Additional Information: full citation, abstract, citings, index terms

Database applications such as enterprise resource planning systems and customer relationship management systems are widely used software systems. Development and testing of database applications is difficult because the program execution depends on the persistent state stored in the database. In this paper we show that how versioning of the persistent data stored in the database can solve some critical problems in the development and testing of database applications can be solved by vers ...

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next